ATLAS Feedback and Priorities for HEP-CCE Year 2+

Torre Wenaus (BNL) for US ATLAS

HEP-CCE All-Hands Meeting

November 6 2020

Outline

- HL-LHC computing requirements: main messages
- HL-LHC computing R&D in US ATLAS, and where HEP-CCE fits in
- HEP-CCE Year 1 appraisal and Year 2 priorities as seen by US ATLAS

HL-LHC CPU Requirements

Meeting the requirements demands aggressive R&D succeeding and reaching production

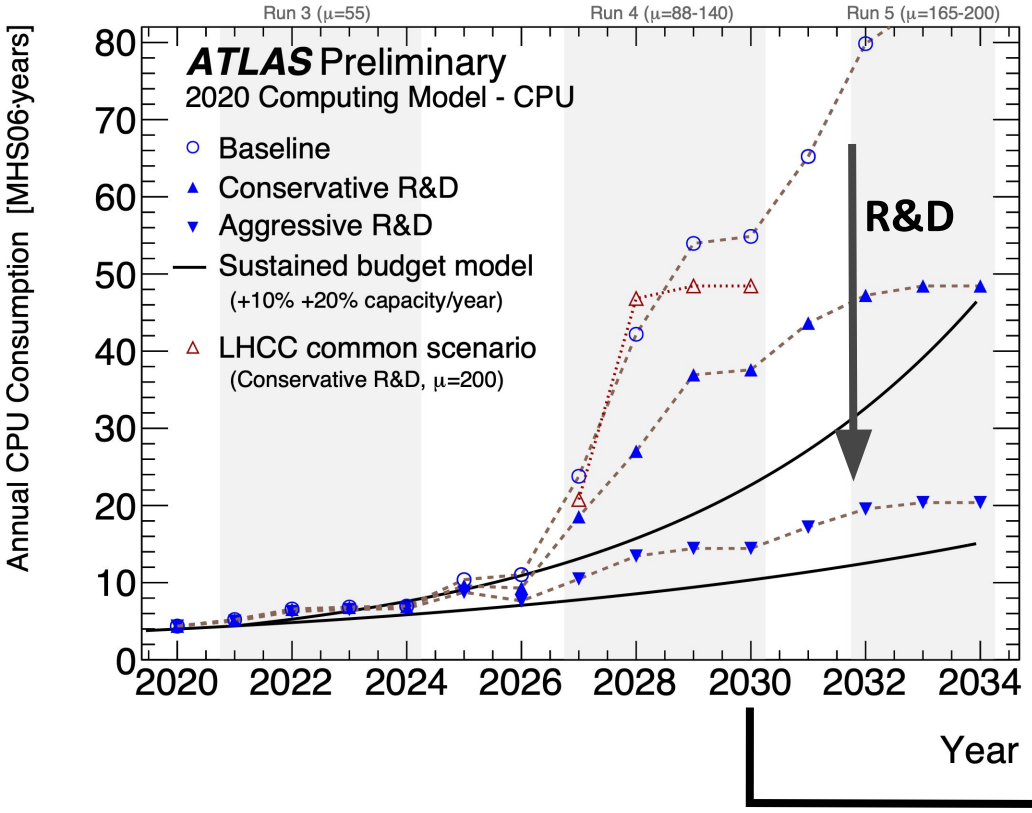

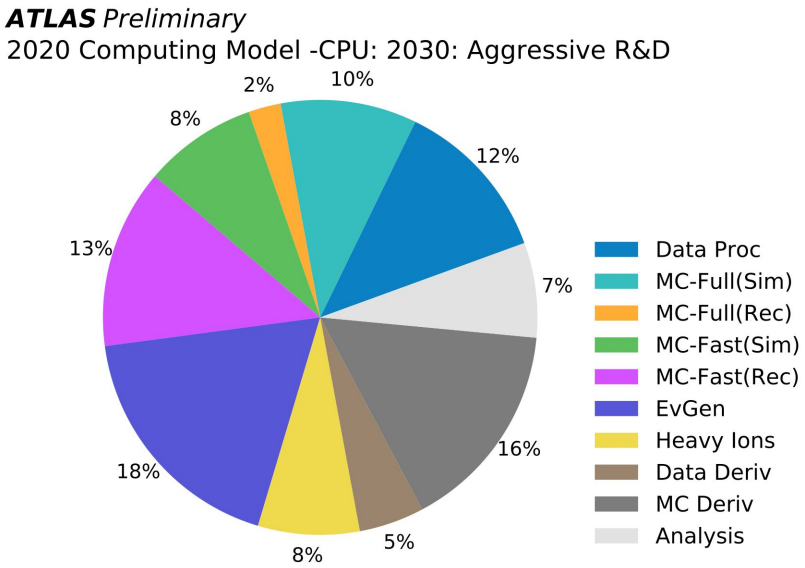

CPU usage is distributed across a wide range of workflows

HL-LHC Disk Requirements

Meeting the requirements demands aggressive R&D succeeding and reaching production

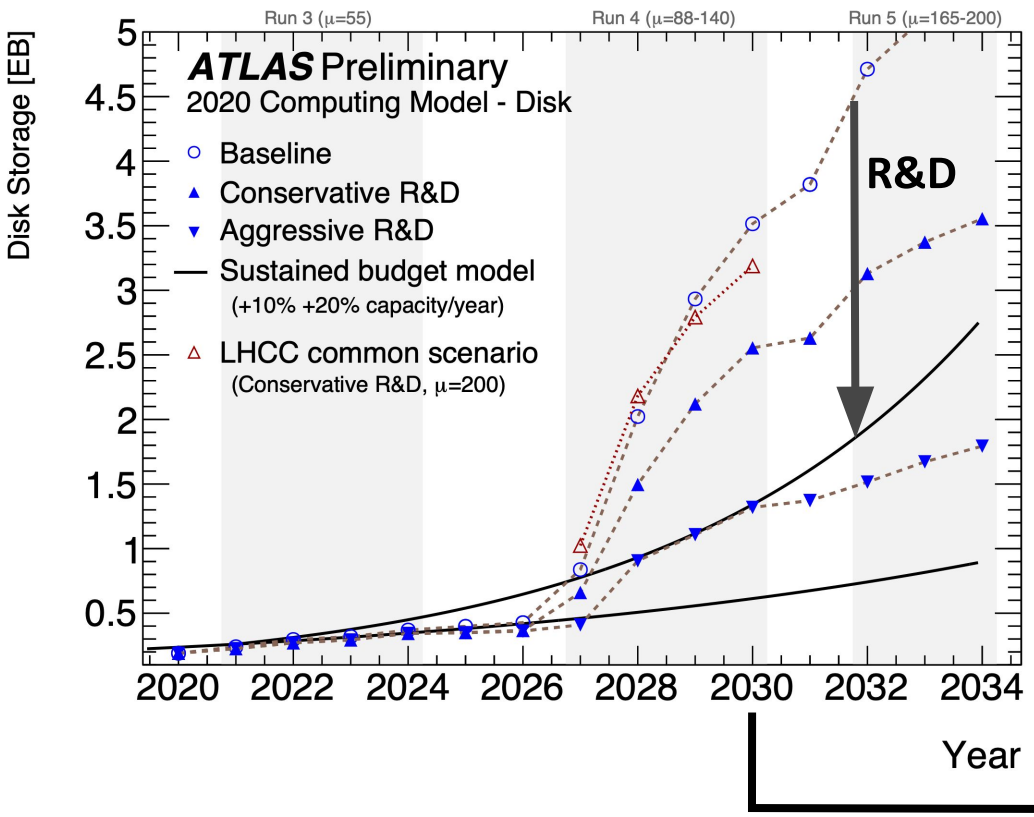

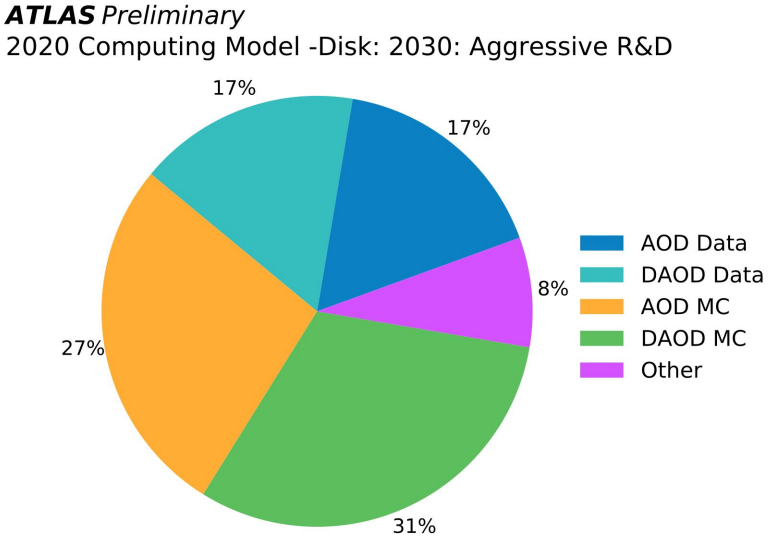

Disk is heavily dominated by analysis formats

HL-LHC Computing R&D in US ATLAS

Created HL-LHC Computing activity (WBS 2.4) in June 2018

Three focus areas chosen for:

- Impact on US ATLAS analysis,
- NSF/DOE priorities (HPC),
- US expertise and long-term commitments

Effort fully integrated with other ops program areas and wider ATLAS activity, and with HEP-CCE, IRIS-HEP, HSF and other community projects

Six of our milestones draw directly on HEP-CCE and IRIS-HEP deliverables

Software re-engineering and algorithm development (2.4.1)

7 FTEs FY21

Workflow porting to new platforms, expected to grow in the future (2.4.2)

0.8 FTEs FY21

Distributed computing (2.4.3)

5.6 FTEs FY21

2.4.1 Software Re-engineering

❖2.4.1.2 Event generation. FTEs: DOE+NSF 0.5, External 0.3 (ANL, SciDAC)

- **Help catalyze the generator developer community to deliver next-gen LCF capable generators.**
- Small activity, hire recently made and ramping up.

❖2.4.1.3 Simulation. FTEs: DOE+NSF 1.8

CCE PPS

- **Commission fast simulation techniques to help solve the HL-LHC CPU problem. FCS would give us an order of magnitude; FC an additional two orders of magnitude.**
- Major activity focused on fast calorimeter simulation (FCS), fast chain (FC).

❖2.4.1.4 Reconstruction. FTEs: DOE 0.3, NSF 0.5, External 0.8 (IRIS-HEP)

- **Help make ACTS the Run-4 tracking software, with needed performance and leveraging accelerators.**
- New activity in FY21 catalyzed by IRIS-HEP involvement in ACTS.

❖2.4.1.5 Analysis

- No in-house ops program activity. Will leverage IRIS-HEP work.

❖2.4.1.6 Framework and services. FTEs: DOE+NSF 2.55, External 0.8 (HEP-CCE)

CCE PPS

- **Core software support for concurrency, accelerators and ML, directed particularly at utilizing accelerators and LCFs in Run-4.**
- Portable parallelization, Raythena on HPCs, next-gen task scheduler, GPU and ML support in Athena.

❖2.4.1.7 Event I/O and persistency. FTEs: DOE 1.7, External 0.5 (HEP-CCE)

CCE IOS

- **Event persistency for Run-4 supporting all use cases and particularly HPCs, accelerators and fine grained workflows.**
- Major activity to support ATLAS event persistency in Run-4.

2.4.2 New Platforms

❖2.4.2.1 Common Platform Infrastructure. FTEs: DOE 0.5 (ANL, UTA)

- **Develop common platform for grid, HPC, and cloud resources**
- Performance measuring and monitoring, release containerization

❖2.4.2.2 HPC and Exascale platforms. FTEs: DOE 0.25 (ANL)

CCE PPS

CCE IOS

- **HPC and exascale implementation and porting of ATLAS production workflows**
- No direct activity until viable LCF applications are developed. Once they are, CCE-PPS and CCE-IOS should both ultimately contribute.

❖2.4.2.3 Cloud

- **Cloud (including commercial) porting and integration**
- Current work is through our collaboration with Google (under 2.4.3)

2.4.3 Distributed Computing

- 2.4.3.2 Data management. FTEs: DOE+NSF 2.35, External 0.5 (DOE Computing), 0.5 (IRIS-HEP)
  - **Leading contributions to data management innovations to solve the HL-LHC storage challenge, particularly through dataflow/workflow orchestration and fine grained dataflows.**
  - Major activity in cold/hot storage optimization, distributed fine grained dataflows, dataflow/workflow orchestration, intelligent data distribution, caching.
- 2.4.3.3 Workflow management. FTEs: DOE+NSF 2.0, External 0.6 (HEP-CCE). *Missing effort: 1-2 FTEs LCF/ECP expert assistance with LCF migration.*
  - **Leading contributions to workflow and workload management to solve the HL-LHC CPU challenge, leveraging all possible resources and architectures, and supporting all ATLAS workflows.**
  - Major activity in workflow and workload management, fine grained workflows, intelligent brokering and scheduling, new workflows (e.g. ML, active learning), HPC workflows.
- 2.4.3.4 Analysis services
  - No ops program work at this time apart from the workflow management effort, see function-as-a-service milestone in 2.4.3.3.
- 2.4.3.5 Common infrastructure. FTEs: DOE 0.75
  - Information services, performance metrics, monitoring, containerization. Infrastructure in support of the other activities.

HEP-CCE, IRIS-HEP Partnerships

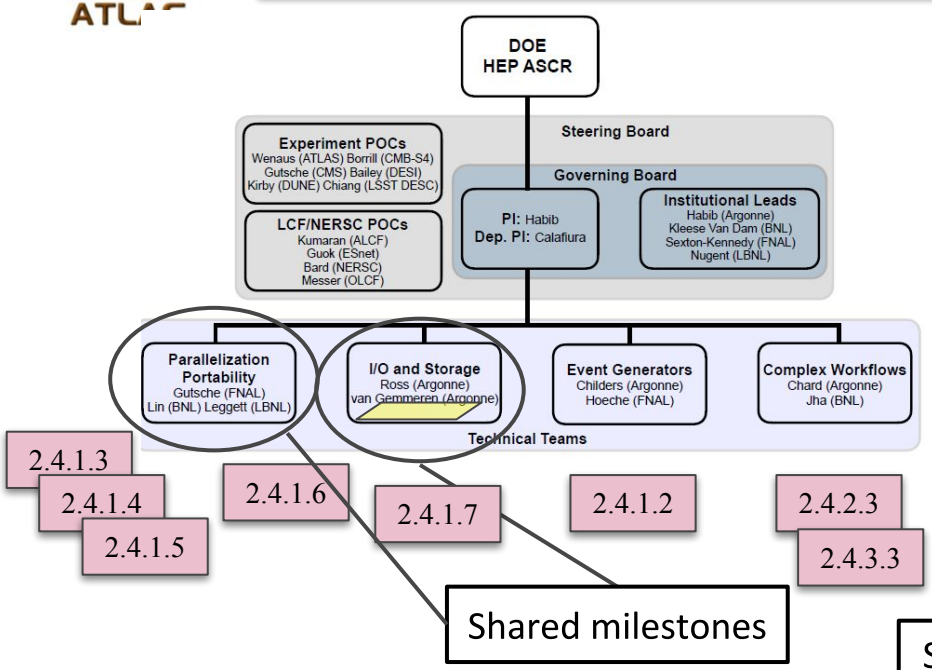

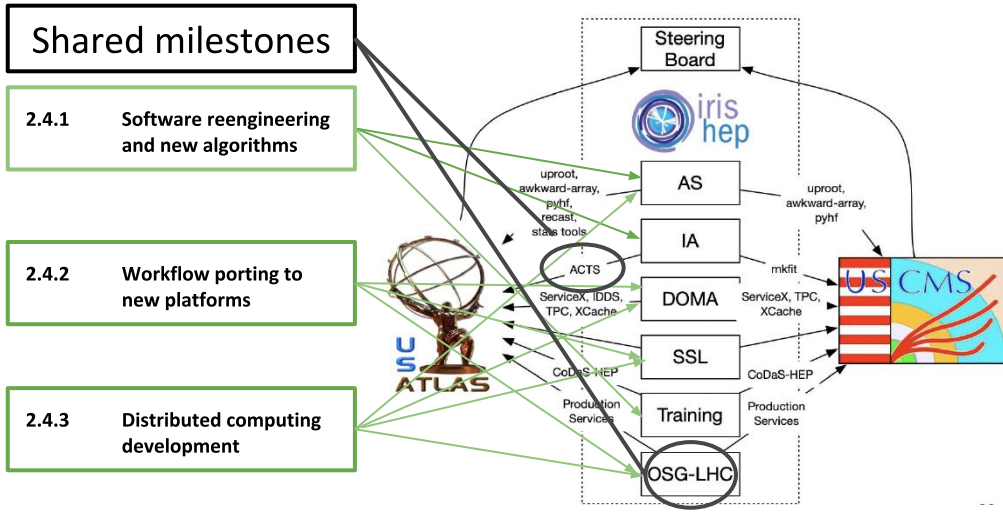

U.S. ATLAS HL-LHC S&C Milestones

Colored: official shared milestone
No color: CCE effort contributes
Grey: Year 2?

	Date	Milestone	Activity
	Oct 2019	FastCaloSim first version working on LCF. Completed.	Simulation
	Dec 2019	PanDA based distributed training and hyperparameter optimization service operating on grid. Completed.	Data & workflow management
	Sep 2020	PanDA based distributed training and hyperparameter optimization service operating on an LCF HPC. Delayed beyond Sep 2020 due to LCF complexity. Would be aided by direct OLCF/ALCF expert assistance.	Data & workflow management
	Jan 2021	Deploy Raythena for running Event Service simulation production.	Core software
	Jan 2021	Develop an automated procedure and software tools for ATLAS software release containerization and validation.	Software infrastructure
	Feb 2021	PanDA/iDDS based ML hyperparameter optimization and training services in production, including user client tools and centralized visualization of training results.	Data & workflow management
	Jun 2021	PanDA/iDDS support for directed acyclic graph (DAG) based workflows to support complex diverse experiment-agnostic workflows including active learning.	Data & workflow management
CCE PPS	Jul 2021	FastChain first version on LCF. Delayed from Jan 2020 due to LCF complexity.	Simulation
	Jul 2021	Complete integration of the new Sci-token identity management system into XRootD.	Data management
	Aug 2021	Evaluation of GPU resource management techniques and infrastructure, with integration into Athena.	Core software
CCE PPS	Sep 2021	Demonstrate the physics and technical performance of the ACTS reconstruction algorithms for ATLAS.	Reconstruction
	Sep 2021	Develop a prototype of the HPX-based task scheduler in Gaudi.	Core software
	Sep 2021	PanDA/iDDS Ray integration supporting checkpoint-capable ML hyperparameter optimization and training on HPCs and other large preemptible resources.	Data & workflow management
CCE PPS + IOS	Oct 2021	Develop a prototype Ray-based mechanism for intra-node scheduling of Athena reconstruction components on heterogeneous platforms.	Core software
	Oct 2021	PanDA/iDDS based function-as-a-service platform allowing users to run fitting and active learning workflows on remote GPU, TPU, and HPC resources. Effectively a joint milestone with the analysis activity area.	Data & workflow management, analysis
CCE IOS	Oct 2021	Complete I/O framework capable of handling EventService Simulation in AthenaMT mode, with correct handling of in-file metadata.	Event I/O
	Dec 2021	iDDS - Data Carousel integration complete for Run-3 production, with iDDS based fine grained orchestration of JEDI, PanDA, and Rucio to minimize task execution time tails.	Data management
	Dec 2021	US ATLAS R&D assessment. All R&D activities deliver a report on their achievements and the steps leading to the 2022 demonstrator. Plans are assessed and ranked by priority and risk.	Planning
CCE PPS	Mar 2022	Deliver a portable parallelization package recommendation to ATLAS software for migration to heterogeneous architectures.	Core software
CCE IOS	Mar 2022	Prototype implementation of non-ROOT parallel event storage backend for fine-grained processing. First R&D target is HDF5.	Event I/O
CCE IOS	Oct 2022	Implement I/O roadmap recommendations to improve metadata support for fine-grained workflows.	Event I/O
	Dec 2022	Deliver Run-4 R&D proof of concept demonstrators for US R&D projects.	Planning
	Dec 2023	Technical design report completed.	Planning
	Jun 2024	Deliver functionally complete prototypes for Run-4 directed US development projects.	Planning
	Jun 2025	Development completion plan in place for Run-4 directed US development projects.	Planning
	Oct 2026	Run 4 dress rehearsal. Test software and distributed computing readiness for datataking.	Planning

HEP-CCE Related Milestones

Zooming in on the milestones related to HEP-CCE...

Date	Milestone	Activity	HEP-CCE
Jul 2021	FastChain first version on LCF. Delayed from Jan 2020 due to LCF complexity.	Simulation	PPS contributing
Sep 2021	Demonstrate the physics and technical performance of the ACTS reconstruction algorithms for ATLAS.	Reconstruction	PPS Year 2?
Oct 2021	Develop a prototype Ray-based mechanism for intra-node scheduling of Athena reconstruction components on heterogeneous platforms.	Core software	PPS + IOS Year 2? This MS relates to the proposed new work on efficient utilization of GPUs
Oct 2021	Complete I/O framework capable of handling EventService Simulation in AthenaMT mode, with correct handling of in-file metadata.	Event I/O	IOS contributing
Mar 2022	Deliver a portable parallelization package recommendation to ATLAS software for migration to heterogeneous architectures.	Core software	PPS joint milestone
Mar 2022	Prototype implementation of non-ROOT parallel event storage backend for fine-grained processing. First R&D target is HDF5.	Event I/O	IOS joint milestone
Oct 2022	Implement I/O roadmap recommendations to improve metadata support for fine-grained workflows.	Event I/O	IOS contributing

HEP-CCE Year 1

- Of the four planned focus areas, the two activated ones PPS and IOS were the right choices for ATLAS
- PPS activities, carrying on from the FastCaloSim GPU study that was already underway, is contributing directly to ATLAS understanding of GPU potential and challenges in the context of a critical real-world app
  - Fruitful collaboration with CS experts drawing on GPU expertise
- IOS activities are addressing the under-explored but critical territory of I/O performance and technologies for high concurrency HPCs
  - Addressing I/O issues in fine grained processing is a direct benefit to ATLAS for its production workflows
  - Quantitative analysis of I/O performance on HPCs drawing on CS expertise and tools is addressing big open questions for ATLAS

HEP-CCE Year 2: PPS

- FastCaloSimV2, used in the PPS studies, is close to production, albeit not in its GPU version
  - Future path has many possible directions, including hybrids between them, still very much R&D
    - Multiple approaches to incorporating ML viewed as most promising
    - GPU R&D continues, branching to exploring new technologies as they emerge and mature
- FCS GPU R&D is relatively advanced in PPS; while it should continue, ATLAS would benefit from broadening the work
- In particular, adding ACTS tracking as an active testbed
  - ATLAS has committed to ACTS: partial use in Run-3, full use in Run-4
  - ACTS has been designed for MT and high concurrency from the start, and has an active parallel ACTS community including IRIS-HEP, with multiple GPU implementations of ACTS algs already being explored, a strong groundwork for understanding the infrastructure issues towards a GPU-capable ACTS.
  - PPS involvement would complement existing parallel ACTS effort in (US) ATLAS, with a particular focus on HPCs and portability
- Also, it would be timely to begin to address the great challenges of achieving significant net throughput benefit, and accelerator utilization, from the accelerator-leveraging applications that do emerge (both accelerator capable apps and ML inference)
  - Bulk (sub)event processing, refactoring framework workflows and scheduling for bulk processing, data marshalling and delivery for bulk processing, Accelerators as a Service to aggregate accelerator usage and achieve high utilization (but without throughput-injuring latencies), ...

HEP-CCE Year 2: IOS

- IOS activities are younger than PPS, much remains to be learned from the Year 1 programs, they should continue
  - Refine understanding of the role of intermediate formats
  - Role of HPC based object stores
- A new activity on efficient accelerator-leveraging workflows would naturally involve IOS as well as PPS activity
- Initial focus on simu is important and well chosen, but attention to analysis data model, formats workflows will be important
  - Analysis data formats dominate disk, today and even more so at HL-LHC
  - Workflows working with these formats (reco, derivations, analysis) represent a lot of HL-LHC CPU, making usability across platforms (including HPCs) important
  - And essential if we aspire to 'virtual data': leveraging all possible processing resources to economize storage by producing-on-demand
  - Challenges are substantial: high IO intensity, moving/using only the subset of data needed, efficient representations for consumers like accelerators and ML
    - Next-gen systems will be very ML friendly.

HEP-CCE: Beyond PPS and IOS

- A year in, PPS and IOS still look to us like the right choices
- Event generators
  - US ATLAS is involved in event generator work, at a small level, where opportunities present themselves
    - SciDAC/base supported MADGRAPH GPU efforts of Taylor, Walter et al
    - A new 50% US ATLAS postdoc working with the orbit of that ANL centric activity
    - CCE could endorse this activity
  - If event generator opportunities present themselves, we should consider them, be opportunistic, but the particular requirements of collaboration with theorist authors, readiness of the software for HPC leverage, and availability of expertise limit the possibilities
- Complex workflows
  - The program described in the proposal is focused on cosmology and on technologies not relevant for ATLAS complex workflows R&D
  - A new activity on efficient accelerator-leveraging workflows would have complex workflow aspects of course, but coupled very tightly within PPS and IOS, we don't need a new activity area

Conclusions

- HEP-CCE Year 1 has been a success
- In the well chosen activity areas of PPS and IOS
- Work has contributed materially both to near term ATLAS priorities and long range R&D towards understanding using accelerators and HPCs
- Both PPS and IOS have productive ongoing programs into Year 2
- Both could benefit from broadening, we've suggested some areas
- Strengthening PPS and IOS looks like the most beneficial approach to us, for Year 2 at least
- We propose that HEP-CCE buys into or endorses the extended list of potential joint milestones we described

Supplemental

U.S. ATLAS HL-LHC Computing WBS